AIAA 2003-0648

Response Surface Methods For Spatially-Resolved Optical Measurement Techniques (Invited)

P. M. Danehy, A. A. Dorrington†, A. D. Cutler#, R. DeLoach
NASA Langley Research Center
Hampton, VA, USA

†National Research Council
#George Washington University

41st Aerospace Sciences Meeting & Exhibit
6–9 January 2003
Reno, Nevada

RESPONSE SURFACE METHODS FOR SPATIALLY-RESOLVED OPTICAL MEASUREMENT TECHNIQUES (INVITED)

P. M. Danehy*
Instrumentation Systems Development Branch
MS 236, NASA Langley Research Center
Hampton, VA 23681-2199

A.A. Dorrington†
National Research Council Associate
MS 236, NASA Langley Research Center
Hampton, VA 23681-2199

A. D. Cutler#
The George Washington University
MS 335, NASA Langley Research Center
Hampton, VA 23681-2199

R. DeLoach**
Instrumentation Systems Development Branch
MS 236, NASA Langley Research Center
Hampton, VA 23681-2199

Abstract

Response surface methods (or methodology), *RSM*, have been applied to improve data quality for two vastly different spatially-resolved optical measurement techniques. In the first application, modern design of experiments (*MDOE*) methods, including *RSM*, are employed to map the temperature field in a direct-connect supersonic combustion test facility at NASA Langley Research Center. The laser-based measurement technique known as coherent anti-Stokes Raman spectroscopy (*CARS*) is used to measure temperature at various locations in the combustor. *RSM* is then used to develop temperature maps of the flow. Even though the temperature fluctuations at a single point in the flowfield have a standard deviation on the order of 300 K, *RSM* provides analytic fits to the data having 95% confidence interval half width uncertainties in the fit as low as ±30 K. Methods of optimizing future *CARS* experiments are explored. The second application of *RSM* is to quantify the shape of a 5-meter diameter, ultra-lightweight, inflatable space antenna at NASA Langley Research Center. Photogrammetry is used to simultaneously measure the shape of the antenna at approximately 500 discrete spatial locations. *RSM* allows an analytic model to be developed that describes the shape of the majority of the antenna with an uncertainty of 0.4 mm, with 95% confidence. This model would allow a quantitative comparison between the actual shape of the antenna and the original design shape. Accurately determining this shape also allows confident interpolation between the measured points. Such a model could, for example, be used for ray tracing of radio-frequency waves up to 95 GHz, to predict the performance of the antenna.

Nomenclature

F	ratio of selected mean square values
K	Kelvin
n	number of data points used to fit a response surface model
p	number of parameters in a model, including intercept
R^2	ratio of explained to total sum of squares
x	streamwise axis in duct
y	vertical axis in duct
z	spanwise axis in duct
σ	standard error in the regression
ANOVA	analysis of variance

* Research Scientist, Member AIAA
† National Research Council Post-doctoral Associate
** Senior Research Scientist, Member AIAA
# Associate Professor, Senior Member AIAA

CARS	coherent anti-Stokes Raman spectroscopy
$CIHW_{95\%}$	95% confidence interval half width
LOF	lack of fit
MDOE	modern design of experiments
PIHW	prediction interval half width
RSM	Response Surface Methodology, Response Surface Modeling
confounding	executing an experiment so that the change in response cannot be uniquely attributed to a specific factor
$F_{critical}$	a threshold F statistic indicating minimum statistical significance
LOF F-statistic	ratio of lack of fit to pure error components of unexplained variance
model F-statistic	ratio of model mean square to residual mean square
model mean square	variance that can be explained by model
P-statistic	probability that a corresponding *F-statistic* is not significant
sample	a discrete number of observations
site	specific location in the design space (specific combination of independent variables)

Introduction

Optical measurement techniques can potentially allow remote, non-intrusive, accurate measurements of many parameters of interest in scientific and industrial applications. Examples include temperature, velocity, and concentration measurements in aerospace flows, flames and engines. They also include spatial ranging and shape measurements of structures, along with countless other applications. These measurement systems can be expensive to develop and difficult to operate. Furthermore, acquiring the necessary data can be time consuming – particularly when using single-point measurement techniques such as the temperature measurement technique employed in this paper. Thus, to reduce costs and improve data quality, we are investigating the application of formal experiment design and analysis to optimize the outcome of optical measurement experiments. This paper focuses on a very important aspect of formal experiment design known as *response surface methods [or methodology], RSM.*

A response surface is essentially an analytic curve that is fit to experimental data. These curves can be fit in multiple dimensions of the independent variables, for example multiple spatial dimensions. *RSM* has some unique features compared to statistical analysis performed without surface fitting. First of all, based on some prior knowledge of the test environment, they allow the experimentalist to determine how much data to acquire prior to doing the experiment. Second, they dramatically decrease the quantity of data required to achieve a given uncertainty compared to conventional methods. Third, *RSM* can add physical insights into the problem being studied. Fourth, *RSM* allows data to be interpolated between the measurement points with a high (and known) degree of confidence. Finally, *RSM* provides quantitative estimates of uncertainties in the model predictions. Despite these advantages, formal experiment design and *RSM* have received very limited attention in the optical measurement community. Our recent paper[1] applying these methods to improve temperature measurements in a supersonic combustor appears to be the first application of *RSM* to optimize single-point optical measurement techniques.

Formal experiment design dates to the early part of the 20th century when it was developed for agricultural experimentation by Fisher and coworkers.[2] *RSM* is credited to Box and co-workers who developed the method for optimizing industrial processes beginning in the 1940's.[3] Many of Fisher's and Box's concepts were improved and popularized by Taguchi and others in the 1970's. Taguchi's methods were adopted by many large US corporations and many still use these methods today. Beginning in the mid-1990's formal experiment design began to be used in the aerospace industry. The application of formal experiment design to windtunnel testing has been led by NASA Langley Research Center, where the method is known as *Modern Design of Experiments* (*MDOE*). Reference 1 cites several recent applications of *MDOE* to windtunnel testing and aerospace applications at Langley and elsewhere. Windtunnel tests that have used *MDOE* at Langley have typically experienced cost savings of a factor of 3 to 5, improved measurement accuracy, improved customer satisfaction and avoidance of systematic errors.[4] *MDOE* is gaining broader acceptance in the windtunnel community, though there is a cultural inertia that maintains the status quo in applications where *MDOE* could make significant contributions.

In this paper, we will explore the application of *RSM* to two different optical measurement applications: First, the temperature field is mapped in a supersonic combustor using coherent anti-Stokes Raman spectroscopy (*CARS*). We review the most successful aspects of the analysis performed in Reference 1 and extend these to better understand how these methods can be better applied in future *CARS* and similar experiments. In the second part of the paper, the shape of an ultra-lightweight inflatable space antenna is measured using photogrammetry. *RSM* is used to model the shape of the antenna, allowing additional insights about the performance of the antenna. By demonstrating the potential of applying *RSM* to these two virtually unrelated research topics, we hope to

demonstrate that *RSM* can provide potential gains for many other measurement applications.

Modern Design of Experiments: A Summary

Reference 1 details the method of *MDOE* and describes many of the advantages of applying *MDOE* to map the temperature field in scramjet combustors. This section summarizes that discussion.

Conventional experimentation often occurs without a carefully designed plan. Typically, just prior to acquiring data, the set points (such as the spatial locations to make the measurements, the fuel flow rates, etc.) for the experiment are chosen, often in an ad-hoc fashion. The goal of the experiment is often loosely defined, such as "to characterize the flowfield as well as possible". Data is acquired until either time runs out, money runs out, or it is decided that enough data has been obtained. The data is then analyzed and an error analysis is often performed as an afterthought.

MDOE advocates a different plan. To begin with, a quantitative objective, or objectives are specified and agreed to by the interested parties (customer, experimentalist, etc.) prior to designing and performing the experiment. This objective is often stated as a required uncertainty level to be obtained in a measured quantity. Next, an experiment that meets these objectives is designed. As part of this design, simple computations determine how many data points will be required to meet the objectives. If this number of data points is fewer than the total allowed by the available budget (including facility availability, real-dollar costs, etc.), then resources can be saved. Alternately, additional configurations or parameters could be studied while consuming the remainder of the allotted budget. If the budget does not allow for enough measurement points to meet the objectives then the objectives need to be re-defined. Or it may be decided that the experiment should not be performed at all. Performing these computations prior to acquiring data dramatically increases the probability that the objectives will be met. A second outcome of the design stage of the experiment is a test matrix that specifies the set points where the data should be acquired.

Execution of an *MDOE* experiment differs from the conventional approach as well. The most remarkable difference is that the order of data acquisition is randomized to avoid systematic errors. For example, spatial locations are visited in a random order instead of acquiring data in a continuous, sequential fashion. Reference 1 provides a compelling hypothetical example to illustrate this point: The parabolic temperature profile in a duct is to be measured, but the temperature is slowly increasing with time. When data is acquired sequentially at spatial locations spanning the duct, temperatures measured at later times at the far side of the duct, are systematically higher than at earlier times at the near side of the duct. With conventional experimentation, this would produce a systematic error in the results: a skewed temperature profile. Because of the sequential order in which the data was obtained, the time-varying systematic effect could not be separated from the spatially-varying temperature profile. These two independent variables (spanwise direction and time) are said to be *confounded*. However by randomizing the order in which the data is acquired, these two effects are separated and the true temperature profile in the duct can be recovered. In many cases such as this one, randomization effectively converts systematic errors into a random error. This is good because random errors can then be reduced by a technique known as *replication*.

Replication is the process of repeatedly obtaining measurements at the same set point (same spatial location, for example), increasing the volume of data acquired, which reduces the measurement uncertainty. Under common experimental conditions, the measurement uncertainty scales as $1/\sqrt{n}$ where n is the number of data points acquired. Replication also provides a statistical reference distribution that can be used to judge the precision of the data. This becomes important when curve fitting the data with response surfaces, as described below. Ideally, replicated measurements should not be obtained at the same set point consecutively. Rather, the set point is varied and the original set point is revisited later.

During the experiment or after the completion of the experiment the data is analyzed using *RSM*. *RSM* is a systematic statistical method for modeling data sets using analytic functions. Basically, curve fitting of a response surface is used to understand and represent the experimental data. Several tests are performed on the curve fits to assure that the model adequately represents the data. Many of these tests are performed during the *analysis of variance* (*ANOVA*). *ANOVA* is a method of partitioning the variance of a data set into two parts: the variance explained by the model and that not explained (unexplained) by the model. The ratio of the explained-to-unexplained variance is known as the *Model F-statistic*. This parameter gives an indication of the ability of the model to represent the data; it can be thought of as a signal-to-noise ratio. For the conditions of the experiments in the current paper (thousands of data point in each fit and $p>10$), models having *Model F-statistics* greater than about 20 pass this test.[1] *ANOVA* further partitions the unexplained variance into two parts: *pure-error variance* and *lack-of-fit variance*. The pure-error variance is computed from replicated measurements. This provides an

estimate of the scatter in the data, whether this is from an imprecise measurement instrument or from shot-to-shot fluctuations in the quantity being measured (for example, turbulence). The amount of unexplained variance that cannot be attributed to pure error is attributed to lack of fit (bad fit to the data). The lack-of-fit, or *LOF F-statistic* is the ratio of the lack-of-fit variance to the pure-error variance. We want to minimize lack of fit and minimize the *LOF F-statistic* to obtain a good fit. Typically, models having a *LOF F-statistic* of unity are acceptable. This indicates that the pure error variance is at least as large as the lack of fit variance, meaning that the model fits the data well. Furthermore, an associated statistic, the *LOF P-statistic* indicates the probability that the value of the *LOF F-statistic* was obtained by chance. This statistic should be >0.05 to assure that the model adequately represents the data.

A figure of merit that describes how much of the variance in the data set is explained by the model is called R^2. If R^2 equals 0.8, then the model explains 80% of the variance in the data. While there is no threshold value for an acceptable R^2, it does provide useful information for model-to-model comparison.

If a model passes the *Model F-statistic, LOF F-statistic,* and *LOF P-statistic* tests, then graphs of the residuals (the difference between the model prediction and the data) should be examined in search of trends. Residuals should be plotted versus time, versus independent and dependent variables, etc. A model that fits the data well will show no obvious trends in the residuals plotted against any of these variables. In that case, the variability of the residuals is due to random chance only. If systematic trends in the residuals are observed then additional terms can be added to the model to quantify and account for these trends.

If the model passes all the tests and there are no significant trends in the residuals, then we can estimate the uncertainty in the model from the approximate relation:

$$CIHW_{95\%} = 2\sigma\sqrt{\frac{p}{n_{total}}} \qquad (1)$$

where $CIHW_{95\%}$ is the 95% confidence-interval-half-width uncertainty in the model, σ is the fit standard error (standard deviation of the residuals), p is the number of parameters in the model, and n_{total} is the total number of data points used in fitting the model. Note that this equation computes a $CIHW_{95\%}$ that is averaged over the range of the independent variables. In fact, the $CIHW_{95\%}$ is typically smaller near the center of the range, and higher near the edges where there is less data.

If the uncertainty computed from Eq. (1) fails to meet the precision requirements specified in the objective of the experiment, then either a better model for the data must be found, or more data must be acquired. On the other hand, if the uncertainty computed from Eq. (1) meets or exceeds the precision requirements, then the model is a candidate for testing. The model is tested against data obtained at the same time as the rest of the data but withheld from the model fit. Typically, about 5% of the data are withheld from the fitting procedure to assess model adequacy. For the model to be accepted, it must successfully predict these withheld data. A successful prediction is defined as a data point that falls within the model's 95% confidence prediction interval half widths ($PIHW_{95\%}$), which are equal to the value of the model prediction $\pm 2\sigma$. As detailed in Reference 1, for 10 trials, 7 or more success are required of a model to achieve a 99% probability that the 95% prediction interval of an adequate model is being verified. For 100 trials, 89 or more successes are required. If the model passes this final test then it can be applied with a known level of confidence for a variety of uses.

Potential Gains from using *RSM*

Potentially the most important benefit of applying *RSM* to optical measurement techniques is an improvement in the precision of the final model. Conventionally, repeated measurements are performed at a single spatial location to determine an accurate mean value. The uncertainty in this mean decreases as $1/\sqrt{n_{point}}$ where n_{point} is the number of measurements obtained at that point. A marked improvement in precision can be obtained by using *RSM* to fit a curve through data measured at multiple spatial locations. One might question how simply fitting a curve through data could possibly decrease uncertainty in the mean at a point. It appears too good to be true. The key idea is that by fitting a curve through the data information is shared from neighboring spatial locations. Suppose temperature measurements have been obtained at five locations along a line where the temperature is slowly increasing. If another spatial point is desired between two of the existing points, it would not take many measurements to determine whether the result at the new location agrees with the trend established by the other points. Information is shared and fewer measurements are required to achieve the same uncertainty. Using the conventional single-point averaging method, the experimentalist starts "from scratch" at each new spatial location and many measurements are required to reduce the uncertainty in the mean.

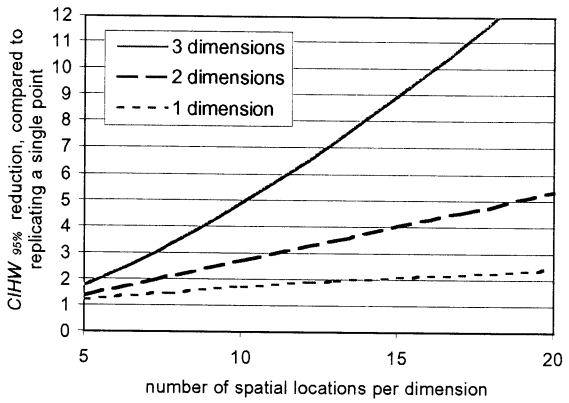

Fig. 1. Factor of improvement in uncertainty by using *RSM* compared to single-location means, plotted against density of measurement points, assumed to be on a rectangular grid. A 6th-order fit in each dimension. Half of its terms in the 6th-order fit are assumed to be significant.

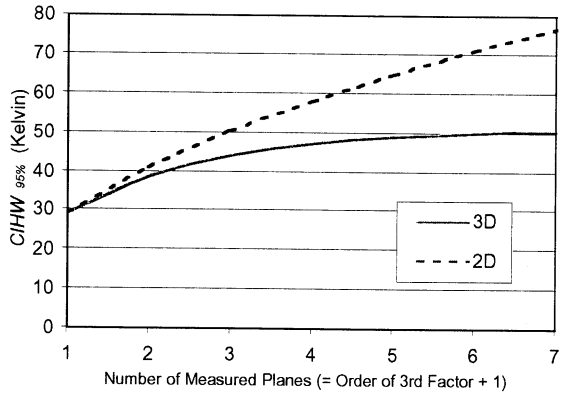

Fig. 2. Uncertainty in fit plotted against the number of measurement planes for (multiple) 2D and (a single) 3D response surface models. For this plot, σ = 300 K, n_{total} = 6000, which corresponds to ~2 days of tunnel operation.

Equation (1) allows us to quantify this potential increase in precision using *RSM*. Figure 8, in Reference 1 uses Eq. (1) to show that when a large number of spatial locations are measured, *RSM* provides lower uncertainties compared to the conventional method of averaging at a single-point. The improvement in measurement uncertainty depends, however, on the number of parameters required in the response surface model, which depends upon on the curvature of the surface. For measurements spread out over 100 spatial locations (a 10 x 10 grid, for example), a 3-parameter response surface would provide nearly 6 times smaller $CIHW_{95\%}$ uncertainties compared to conventional single-point means, assuming the same number of total data points were acquired. A 10-parameter surface would have $CIHW_{95\%}$ 3 times smaller than single-point means, and a 30-parameter surface would have $CIHW_{95\%}$ nearly 2 times smaller than single-point means. These improvements in $CIHW_{95\%}$ relate directly to the measurement volume required since $CIHW_{95\%} \sim 1/\sqrt{n_{total}}$. Thus, to obtain similar $CIHW_{95\%}$ with the two approaches, *RSM* would require 1/30th, 1/10th, and 1/3rd the amount of data for p = 3, 10, and 30, respectively, compared to conventional methods.

Reference 1 concluded that using *RSM* in the analysis resulted in a 5-fold reduction in required data volume compared to single-point means, for fitting single measurement planes of data. One of the goals of the present paper is to extend and optimize this approach for the present experiment to further increase measurement efficiency and gain further knowledge about the flowfield. In particular, we wish to investigate the possibility of spreading measurement points across multiple planes of data, either spatially or to investigate how a single spatial plane changes while adjusting another independent variable, such as fuel equivalence ratio.

Suppose temperature measurements were to be made in a 3-dimensional volume. Figure 1 shows the improvements in $CIHW_{95\%}$ that would be expected compared to measuring and computing single-point means. The figure was computed assuming that 6th-order polynomial models would be fit to the data in each dimension. Furthermore, it is assumed that only half of the total number of parameters in these models would be significant, as is frequently the case. Finally, it is assumed for the purposes of this computation, that measured surfaces are smoothly varying, though this is not a requirement for using *RSM*. In one spatial dimension, the gains are modest: fitting a 6th-order model over 10 spatial locations yields uncertainties 1.7 times smaller than single-point means. However, using a 6th order response surface to fit to a measured over a plane (eg. a 10 x 10 grid) provides a 2.5 reduction in uncertainty. Finally, fitting a 6th-order model in each of 3 dimensions over a 10 x 10 x 10 grid provides a 5 times reduction in uncertainty. Stated another way, it would take only 1/25th the amount of data to characterize the 3D measurement volume with *RSM* compared to computing single-point statistics, assuming equivalent uncertainties. This is a *tremendous* cost savings. An added benefit of fitting the data with a response surface is that data can be reliably interpolated between measurement points if the quality of the fit is good. Of course, fitting over 4 dimensions (*e.g.* 3 spatial dimensions and fuel-equivalence ratio) provides even bigger cost savings compared to conventional methods.

In the recent *CARS* thermometry experiment reported in Reference 1, temperature was mapped at five different streamwise locations. Based on the predictions of Figure 1, it would seem prudent to fit a single three-dimensional model through all five planes.

What improvement in $CIHW_{95\%}$ would be expected for this situation? Figure 2 addresses this question. To compute this figure, we assumed n_{total} = 6000 and that σ = 300 K. For this figure, 6th-order models were assumed in each of the two dimensions in the measurement planes, and the fit order for the third dimension is taken to be equal to the number of planes fitted minus one. Thus, if two planes were measured, a linear model would be fit in the third dimension; three planes would allow a quadratic model to be fit in the third dimension and so on. Half of the terms in the model were assumed to be significant. Consequently, models for fitting 1, 2, 3, 4 and 5 planes of data would have p = 14, 25, 32, 37 and 40 terms respectively. The figure shows a 1.3 reduction in $CIHW_{95\%}$ by fitting the 3D volume to 5 planes, compared to fitting individual temperature maps to each plane – a worthwhile improvement, though not as stunning as the trends predicted by Figure 1. Fitting just two planes of data results in a very modest improvement in $CIHW_{95\%}$, 1.07, compared to fitting individual 2D planes. The reduction in uncertainty predicted by this figure occurs for the following reason. When fitting multiple planes of data, n_{total} scales as the number of planes, whereas p grows more slowly as the number of planes is increased. When fitting two planes together instead of separately for example, n_{total} doubles, but p also nearly doubles leaving the $CIHW_{95\%}$ almost unchanged. However, when fitting four planes compared to two, n_{total} doubles again while p increases by less than half, thereby reducing the $CIHW_{95\%}$ substantially.

Even if the fit uncertainty is not improved much, it may still be worthwhile to fit data in 3D to allow interpolation between measurement planes. Another observation from Figure 2 is that for the same $CIHW_{95\%}$ (equal to 50K in the figure) and the same total number of data points, 7 measurement planes could be studied with the 3D model compared to just 3 with 2D models. Thus, higher-structure flowfields could be investigated for the same cost.

The reasons that Figure 2 does not appear to deliver the dramatic potential illustrated in Figure 1 are twofold. First, Figure 1 shows 1D, 2D and 3D response surface fits compared to *single-point* means. The curves in Figure 2 only compare 2D to 3D response surfaces. Both of these types of models, of course, already have substantial improvements over single-point means; these improvements are omitted from the figure for clarity. Second, the enormous gains in efficiency only begin to occur when the number of spatial measurements per dimension exceeds the order of the fitted polynomial. For example, Figure 1 shows huge gains in efficiency when 15 locations are measured per dimension, while a 6th-order model is fit. In fact, only 7 spatial locations are required to fit a 6th-order model in 1D. Why would one measure more spatial locations than this? Measuring more spatial locations allows visualization of higher-curvature flow structures. If flow structures are encountered that are higher order than anticipated then they can be fit with a higher order model during the analysis, though with lower than anticipated $CIHW_{95\%}$. This approach is more conservative than the conventional approach where scarce resources are spread over few spatial locations in an effort to obtain accurate measurements. Thus, the *RSM* approach helps to measure unanticipated higher-ordered spatial structure. If this precaution is used, and the data volume is sampled at many spatial locations, then *RSM* does, in fact, provide substantial costs savings compared to conventional methods.

Figure 2 prompts a further question: If a fixed number of data points is spread across many spatial locations, is there a practical lower limit to the number of measurements that can be made in a single plane, while still yielding sensible data? This point is discussed below.

CARS Temperature Measurements in a Supersonic Combustor

A. Practical Considerations

The description above outlines the way an ideal experiment would be designed, executed, and analyzed using *MDOE*. Unfortunately, some compromises that prevented the full implementation of *MDOE* were required in the present experiment. In this section, we briefly describe the experimental setup as it relates to the current paper and then discuss how *MDOE* was implemented. Much of the discussion in this section originates from Reference 1 though many details have been omitted for brevity.

Experiments were performed in NASA Langley Research Center's Direct-Connect Supersonic Combustion Test Facility (DCSCTF), which is a vitiated, blowdown wind tunnel. The supersonic combustor model consisted of a short rectangular duct containing a rearward-facing step, after which gaseous hydrogen fuel was injected. The supersonic combustor operated with a steady flow time of 10-20 seconds.

We used planar *BoxCARS* to make single-shot broadband nitrogen *CARS* temperature measurements in the combustor.[5] During the flow time, 100-200 single-shot temperature measurements were obtained. *CARS* spectra, acquired from an intensified CCD linear array attached to a spectrometer, were fit with a library of theoretical curves for a range of temperatures and N_2 concentrations to determine the gas temperature on each pulse of the laser. Thus, a database of x, y, z locations, temperatures, and N_2 concentrations was generated for five measurement planes in the flowfield

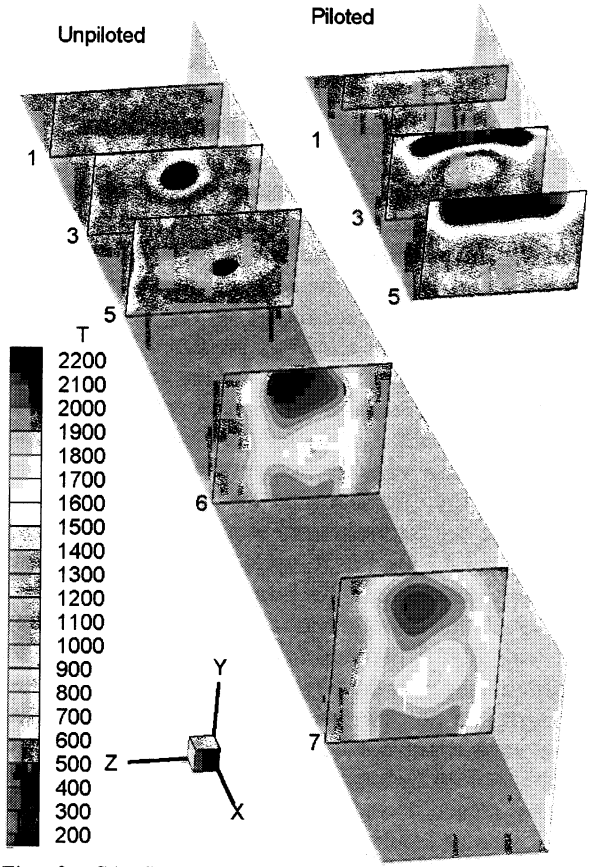

Fig. 3. *CARS* temperature maps in the combustor duct. Temperatures are in Kelvin. See Ref. 5 for details.

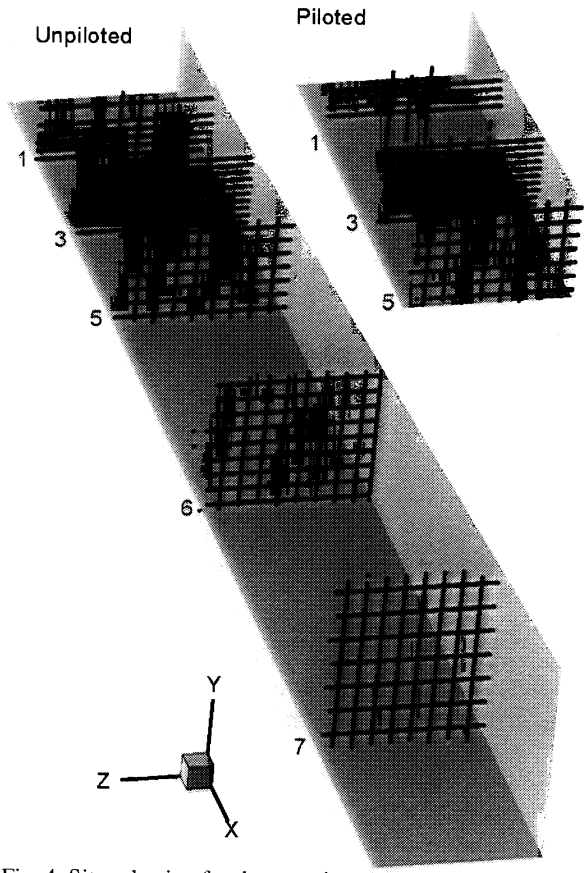

Fig. 4. Site selection for the experiment.

for unpiloted operation and three planes when H_2 pilot fuel was injected upstream of the main fuel injection.

Figure 3 shows the resulting temperature maps obtained in the supersonic combustor.[5] Briefly summarizing the results, the vitiated air flow enters the test section at about 1160 K. For the unpiloted case, cold fuel, with a stagnation temperature of about 300 K, is injected between planes 1 and 3. Evidence of a small amount of combustion occurs around the periphery of the fuel jet in plane 3 where the temperature exceeds the incoming freestream value. The temperature of the cold jet increases between planes 3 and 5, but most of the combustion occurs between planes 5 and 6. Planes 6 and 7 show hot combustion products on the top and bottom of the duct and remnants of the cold fuel jet near the middle of the duct.

Reference 1 described in greater detail how *MDOE* was applied to this experiment and how analytic surfaces were fit to the data, to produce the temperature maps shown in Fig. 3. Reference 1 focused on the analysis of plane 3, unpiloted operation, which was thought to be the most challenging measurement plane to analyze due to the large temperature gradients present. In the current paper, we study other planes, chosen on a case-by-case basis to investigate different aspects of *RSM*.

The objectives for the scramjet thermometry experiment were determined from interviews with the customer, a computational fluid dynamics expert. Together it was decided that the goal of the experiment would be to map the flowfield with a specified precision, in as many planes in the combustor duct as resources would permit. From discussions with the CFD expert, we established ±50 K as the required 95% confidence interval for model predictions. For resource planning purposes, we computed the volume of data necessary to produce a model with such a 95% confidence interval, averaged over all the points used to fit the model. Also, from prior measurements in the same facility we knew that turbulent fluctuations in the flow had a standard deviation of about 300 K. Furthermore, observing the spatial distributions in past measurements, we determined that we would probably fit the data with response surface having approximately 20 parameters. Using Eq. (1) the required number of data points was estimated to be 2880 per measurement plane.

If we had instead used the conventional approach of acquiring many temperature measurements at a single location to reduce uncertainty in the mean, and if

we targeted the same precision requirements, 144 temperature measurements would have been required at each site. The *MDOE* budget of 2880 points per plane then corresponds to only 20 measurement locations if the conventional method had been used. These would perhaps be enough spatial locations to measure the temperature along two or three lines in the flowfield. But, for the same cost, *MDOE* allows us to map the temperature over the entire plane.

Reference 1 discusses the reasons for choosing the spatial measurement locations in detail. To summarize, two types of measurements could be performed at the maximum data rate. One type of measurement was obtained by scanning the probe volume horizontally or vertically through the flowfield using a pair of periscopes driven by stepper motors. A second type was performed at fixed locations in the duct for the entire duration of the run. A combination of these two types of measurements was used to map each measurement plane during a single day of measurements. All fixed locations were distributed uniformly in the vertical direction on the spanwise center of the flow. We alternated randomly between fixed and scanned runs. We also randomized the directions of the scanned measurements (up/down or across/back). Figure 4 shows the sites selected for the different measurement planes in the experiment.

After the experiment was performed, temperatures were computed from the raw *CARS* spectra. See Ref. 5 for details of this procedure. The y, z locations of the measurement volume and the corresponding temperatures from all runs in a given plane were written to data files, one line for each individual temperature measurement.

Two different statistical analysis programs were used to fit the data: Design-Expert®[6] and TableCurve®3D[7]. These programs have many common features, including automatic fitting of surfaces to data and automatic computation of the statistical figures of merit (*Model F-statistic*, R^2, etc). TableCurve®, however, does not allow model terms to be removed as easily as Design-Expert®. Removing model terms is very important in the present application because reducing the number of terms p increases measurement precision, according to Eq. (1). It is often found that more than half of the terms in the models are insignificant. So, by removing these terms and re-fitting the data with the reduced model, one can reduce the uncertainty in the fit by a factor of about 1.4. Unfortunately, the only candidate model functions in Design-Expert® are polynomials, whereas TableCurve® provides thousands of different model forms, including Chebychev, Fourier, Sine and Cosine Series, etc. Thus, both software packages were used in the analysis.

Reference 1 describes how nine different proposed models were evaluated prior to settling on the best one. The current paper uses polynomial fits and also various orders of the Cosine Series Bivariate function.[7]

B. Fitting Data in Multiple Dimensions

During unpiloted operation of the supersonic combustor, temperature was measured at planes 1, 3, 5, 6 and 7. In principle one could fit a 3D response surface model through these 5 measurement planes. Doing so would take advantage of the factor of 1.3 added improvement in $CIHW_{95\%}$ predicted by Figure 2 and also would produce a model for the entire duct that would allow interpolation of the data anywhere in the volume. However, underlying *RSM* is an assumption that the response variable (temperature) is a smoothly varying quantity. If the surface has discontinuities then the data set can be broken up along these boundaries. *RSM* can then be used to fit up to these boundaries and then the solutions can be joined at the boundaries. Because of the smooth nature of the fitted functions used with *RSM*, discontinuities in flows owing to dramatic changes in the physics of the flowfield cannot necessarily be fit. Two such changes in physics occur in the supersonic combustor model. First, a large quantity of cold hydrogen fuel is injected into the flow between planes 1 and 3. Second, after a period of latency, abrupt combustion of the fuel occurs between planes 5 and 6. This is observed by the sudden rise in temperature between these two planes and also is corroborated by pressure measured along the length of the duct. Thus, just two *pairs* of measurement planes: 3 & 5 and 6 & 7 can be fit together sensibly. We analyzed both pairs of measurement planes but will only discuss planes 6 & 7 here.

To fit a single measurement plane, a full 6th-order polynomial model was used in y and z, where y is in the vertical direction and z is spanwise direction (note change in coordinate system compared to reference 1). This polynomial model has 28 coefficients:

$$T = a_0 + a_1 y + a_2 z + a_3 yz + a_4 y^2 + \ldots + a_{27} z^6 \quad (2)$$

Data from planes 6 and 7 were combined into a single Design Expert file. Since there were only two values in the x (streamwise) direction, only models that were linear in x were considered. For the combined measurement planes, a full 6th-order model was fit in y and z. This model, shown in Eq. (3) has a total of 48 coefficients.

$$T = a_0 + a_1 y + a_2 z + a_3 yz + a_4 y^2 + \ldots + b_0 x + b_1 yx + b_2 zx + b_3 yzx + b_4 y^2 x + \ldots + b_{19} z^5 x \quad (3)$$

Equation (3) can be thought of as the same as Eq. (2) but with an additional term linear in x for each term

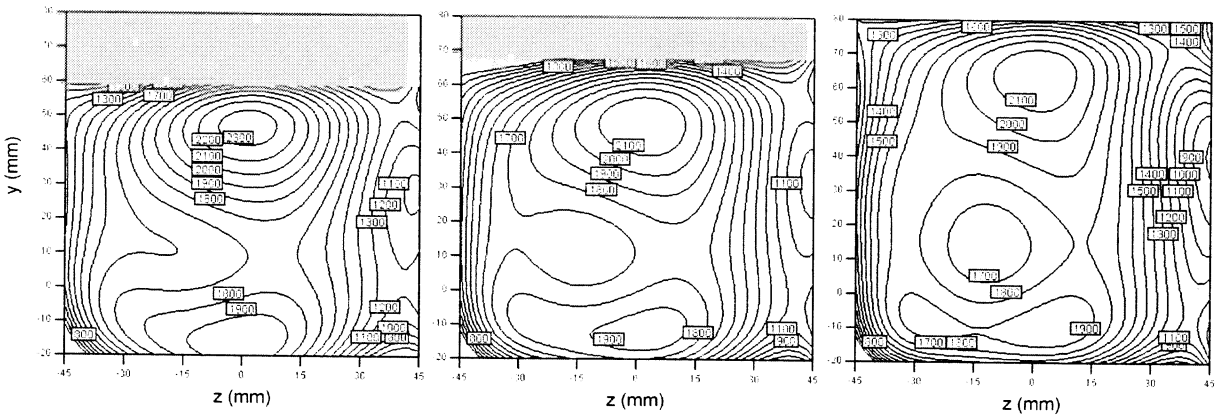

Fig. 5. Temperature maps for Plane 6 (left), Plane 7 (right) and an interpolated plane half way between 6 & 7 (middle).

below 5th-order in Eq. (2). These additional terms effectively allow each coefficient in Eq. (2) to evolve as the flow propagates downstream. Note that the number of parameters, p, has nearly (but not quite) doubled. When considering uncertainty estimates using Eq. (1), the increase in p is offset by the fact that n_{total} is now the sum of the number of points over 2 planes. Thus, at first glance, it appears that the $CIHW_{95\%}$ for fitting two planes will be similar to fitting each plane by itself. Ideally, if the flow changes very little between planes, many of these b coefficients would be insignificant, and could be omitted, thereby reducing p. In the current experiment, for both planes 3 & 5 and 6 & 7 that turned out not to be the case.

Figure 5 shows a contour map of the temperature in planes 6 and 7 and also an interpolated temperature map in the middle. The resulting $CIHW_{95\%}$ for the fit to the pair of planes is 36 K, whereas the $CIHW_{95\%}$ for fitting plane 6 alone was 37 K, and the $CIHW_{95\%}$ for fitting plane 7 alone was 35 K. Thus, even the very small improvement in $CIHW_{95\%}$ of 1.07 predicted by Figure 2 was not realized, though several of the assumptions used to compute Fig. 2 were not exactly satisfied. It is interesting to note that σ for the combined 6 & 7 fit is approximately the average of the values obtained by fitting 6 and 7 separately. This is one indication that fitting the pair of surfaces has not compromised the quality of the fit. Another indication of a good fit to the pair of surfaces is that the R^2 is only slightly worse than the R^2 for the individual plane fits. In summary, it appears that the only advantage of fitting a 3D volume of data in the present example, is that data can be interpolated between the measurement planes.

In these supersonic combustion experiments, the locations of these streamwise measurements are heavily constrained. There are only 7 possible measurement planes. Also, it is nearly impossible to visit more than one streamwise measurement location in a day because the optical table needs to be moved and the *CARS* system must be re-aligned when changing measurement planes. This is not feasible because access to the facility during testing is barred for safety reasons. Thus, it is very difficult to measure the multiple planes of data that allow the gains predicted by Figure 2 to be realized.

Even with the present experimental configuration, there are ways that some of these potential gains can be realized. If the flow of interest had a flame that was stabilized near the fuel injector, then one could argue that planes 3, 5, 6, and 7 (or better yet 2, 3, 4, 5, 6, and 7) experience no fundamental changes in the underlying physics between them. In this case, it would be possible to improve the $CIHW_{95\%}$ by up to a factor of 2 by fitting in 3D, depending on the order of the model required to fit trends in the streamwise direction. This means that those measurement planes could be mapped to a specified level of precision with a volume of data that is 1/4th of the amount that would be required to map each plane independently.

Another way that the present experimental configuration would allow some of the gains predicted by Figure 2 to be realized is to measure a single plane of data as a function of another independent variable, such as fuel flow rate (stoichiometry). Changing stoichiometry does not require access to the wintunnel facility. So, a few different stoichiometries could by studied in a single day: sharing 3000 data points, 1000 at each of at 3 different stoichiometries, for example. In this case, by fitting the data in three dimensions the $CIHW_{95\%}$ would be a factor of 1.5 worse than by obtaining all 3000 points with the same stoichiometry. However, considering that the $CIHW_{95\%}$ obtained in the Reference 1 exceeded the required measurement precision by nearly a factor of two, this is a viable option.

It should be pointed out that many other experiments do not have such severe constraints on the measurement locations because of optical access and facility limitations. In these cases, the true gains predicted by Figures 1 and 2 could be realized.

C. Reducing Number of Measurements Per Plane

The previous section concluded with the suggestion that the total number of data points could be spread over multiple planes, reducing the number of measurement points in each plane while allowing more planes to be studied. A sufficient n_{total} would be acquired to meet the measurement requirements. However, the present experiment is also heavily constrained in the way measurements can be obtained *within* a plane. As described previously, during the 10-20 seconds of a tunnel run, measurements can be either obtained at a fixed spatial location or can be obtained by scanning across the duct at constant velocity. For scanned measurements, the scans must be along straight lines, with one straight line per tunnel run. At issue is the fact that 1000 data points can consist of as few as 5 straight line scans: hardly enough spatial coverage to successfully fit a 6th-order model in two dimensions. Considering the fact that at least one scan must be repeated to generate replicates, that leaves only 4 straight line scans to fit the data over a plane. This simply may not be enough spatial locations to map a plane. So, how many scans is too few? A conservative estimate would be obtained by considering how many scans are necessary to obtain a single plane of data. This section describes a brief empirical analysis to determine approximately the number of linear scans required to recover a typical temperature profile in the supersonic combustor.

To investigate this effect, we fitted data obtained from plane 3 during unpiloted operation. The goal was to fit the data using measurements from fewer and fewer tunnel runs to determine how few scans could be used to develop a good quality model. We fit the data with a consine series bivariate order 6 model. There were data from a total of 19 tunnel runs available for this measurement plane. Fitting all of this data provided a baseline case for comparing with the other scans. The 19-run data set had a $CIHW_{95\%}$ = 32 K. Cutting the number of runs by half, to 10 runs, we obtained a $CIHW_{95\%}$ = 46 K. Removing a few more runs, to a set of just 7 runs, produced $CIHW_{95\%}$ = 54 K. Each of these three temperature maps looked qualitatively similar, and comparisons of the graphs of the fits show that the models agree to within the confidence intervals almost all of the time. However, cutting the number of runs to 5 – a quarter of the original number of runs – resulted in a temperature map that looked very different than the others just described. This final temperature map surely would have failed a test designed to use the model to predict data not included in the fit. Thus, based on this preliminary, empirical study, we believe that we can achieve satisfactory fit quality within the desired confidence level of 50 K, using about 10 scans per plane (for example, a 4 x 4 grid of line-scans with 2 repeated scans). This is half or a third of the number that were acquired in the experiment analyzed in this paper.

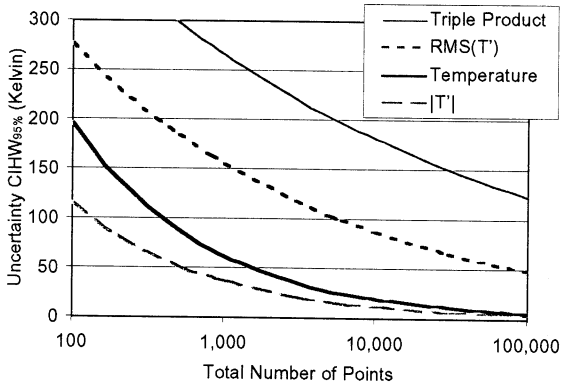

Fig. 6. Predicted uncertainties in mapped quantities, based on four 180-shot single-location temperature measurements in the flow. Measurements were taken along the spanwise centerline in Plane 5, unpiloted operation.

D. Computing Derived Quantities

Everything discussed up to this point has focused on determining accurate *mean* quantities in a supersonic combustor. Frequently, statistical quantities other than means are required. In fluid mechanics and combustion, a particularly important statistical quantity is the root mean square of the temperature fluctuations about the mean, *RMS(T')*, where *T'* is the difference between an individual temperature measurement and the mean at that location. The standard method of computing *RMS(T')* is to compute the mean of a group of measured temperatures at one spatial location, and then to average the squared difference between each sample and the mean. Taking the square root of this number produces the *RMS(T')*. This figure is a useful quantity for several reasons. First of all, it provides a measure of turbulent fluctuations in the flowfield. More highly fluctuating parts of the flow have a larger *RMS(T')*. Secondly, statistics of the fluctuations, such as this, are often used by turbulence models employed in computational fluid dynamics codes.

In Reference 5 we proposed the extension of response surface methods to map the *RMS(T')* distribution across the flowfield. The new method of computing the *RMS(T')* surface uses the model prediction (obtained previously) as an estimate of mean temperature throughout the flow. The squared differences between each data point and the model prediction is then fit using *RSM*. Taking the square root of this fitted surface allowed us to determine the *RMS(T')* map.

Other functions that provide information on turbulent fluctuations in the flow include the absolute value of the temperature fluctuations, $|T'|$, and the cube root of the triple product, defined as $\sqrt[3]{T'^3}$. In principle, all three of these derived quantities can be mapped over a surface using *RSM*. Furthermore, we can determine the number of data points required to reach various levels of confidence with *RSM*. First, the standard deviation, σ, of the quantity being fit must be known. In the present experiment, this was determined by measuring temperature repeatedly at a few typical points in the flowfield. We sampled the temperature 180 times at four different spatial locations in plane 5. From these measurements, we could determine the relevant standard deviations for use in Eq. (1): $\sigma(T)$, $\sigma(|T'|)$, $\sigma(T'^2)$, and $\sigma(T'^3)$. Averaging together values from the four spatial locations gave an estimate of each σ over the measurement plane. These average values were $\sigma(T)$ = 254 K, $\sigma(|T'|)$ = 149 K, $\sigma(T'^2)$ = 98,000 K^2, and $\sigma(T'^3)$ = 77 x 10^6 K^3. We assumed that p was equal to ~15 which is typical for the surfaces fitted in this paper. Finally, n_{total} could be specified to determine the expected uncertainty or alternately, an uncertainty specified to determine the n_{total} required.

Figure 6 shows the comparison between the predicted uncertainties for 4 parameters of interest. Regarding mean temperature maps, the figure predicts that $CIHW_{95\%}$ = 50 K can be achieved with about n_{total} = 1000. Only 500 measurement points are required to obtain $CIHW_{95\%}$ = 50 K for mean maps of $|T'|$. However, to obtain maps of *RMS*(T'), the model predicts that n_{total} = 90,000 is required to obtain $CIHW_{95\%}$= 50 K. Finally, in order to obtain an uncertainty of 50 K in a map of the cube-root of the triple product of T' would require n_{total} = 20 million!

Figure 7 shows temperature maps of plane 5 for three of these quantities, T, $|T'|$, and *RMS*(T'). All these maps are based on the same set of n_{total} = 3111 measured temperatures. For the average temperature map the $CIHW_{95\%}$ = 42 K, in reasonable agreement with the prediction of 35 K from Figure 6. For the map of $|T'|$ fit using the Cosine Series Bivariate Order 6 function, the $CIHW_{95\%}$ = 20 K in close agreement with 21 K predicted from Figure 6. For the *RMS*(T') map, the $CIHW_{95\%}$ = 126 K again in reasonable agreement with 118 K predicted by the Figure. Thus, the model appears to correctly predict the confidence levels for these three parameters. Variation between predicted and actual $CIHW_{95\%}$ is expected because p varies from fit to fit and also because it is possible that the four measurement locations used in the predictions don't represent the whole plane.

A significant result of this exercise is to show that it is not possible to achieve maps of *RMS*(T') or cube-root of the triple product of T' with 50 K uncertainties in the current experiment. For example, obtaining the required 90,000 data points would require about 30 days of tunnel operation to acquire enough data for a single measurement plane! This is greater than the normal amount of tunnel running time budgeted for an

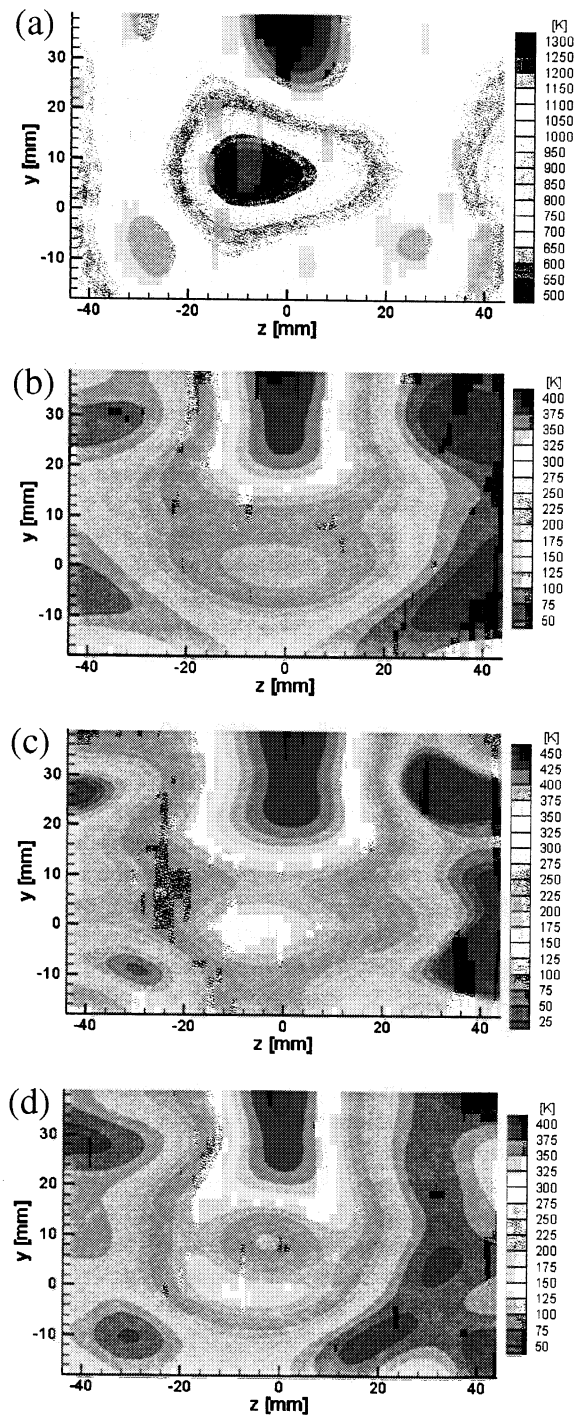

Fig. 7. Cosine Series Bivariate Order 6 surfaces fitted to T (a); $|T'|$ (b); *RMS*(T') (c); and Cosine Series Bivariate Order 8 surface fitted to $|T'|$ (d). Each is for Plane 5, unpiloted operation.

entire experiment. Surely this is not the best use of resources. Thus, if this type of information is desired with a high degree of confidence, it might be better to perform measurements along a line, or perhaps at just a few single points in the flow. For this reason, it might be desired to study a much simpler flow, such as a two-dimensional mixing layer in which there is much less spatial structure, and models having many fewer coefficients, p, can be fit to the data. For example, if a fifth-order model in one spatial dimension was fit to the data obtained through a mixing layer, the $RMS(T')$ profile could be measured with $CIHW_{95\%}$ = 90 K in a single day's run (n_{total} = 3000).

These statistics seem fairly discouraging for the present experiment where the data rate is slow (10 Hz), the tunnel runs are short (20 sec.) and where the wind tunnel can only be run a limited number of times per day (~20-30). However, there are many other types of measurement systems that can be used to study turbulent fluctuations, and the formalism described here can be used to address these experiments as well – hopefully with a better outcome. Nonetheless, it is extremely valuable to know prior to performing the experiment whether the requirements of the experiment can be met.

Panel (d) in Figure 7 shows the Cosine Series Bivariate Order 8 function fit to the $|T'|$ data. This map clearly shows the round fuel jet slightly above the origin (0, 0). The fuel jet is marked on the edges by highly-fluctuating regions, with a more quiescent region in the middle of the fuel jet. Fitting the same data with the slightly lower order model, in panel (b) above, obscured the structure of the fuel jet. This comparison provides a convincing example of why one should over-sample the design space spatially. If a sixth order model is anticipated, it is wise to acquire data at enough spatial locations to fit a seventh or eighth order model should it later be required. On the contrary, for the mean temperature maps, seventh and eighth order fits looked qualitatively similar to the sixth order fit. Since the sixth order fit also passed all the statistical tests, it was chosen because it adequately represents the data while having the fewest parameters, p, resulting in the lowest possible uncertainties.

Photogrammetric Shape Measurement of Space Antenna/Solar Collector

A. Test structure and measurement procedure

A study was undertaken to evaluate the potential of *RSM* to develop a numerical model describing the shape of a physical object, in this case a 5-m inflatable, parabolic space antenna, shown in Figure 8. This ultra-lightweight space structure was manufactured for NASA in 1996 by SRS Technologies in Huntsville, Alabama. It was chosen for this study because it was thought that its parabolic shape would lend itself favorably to a polynomial surface model, and that the coefficients of that model could provide useful information about the performance of the antenna.

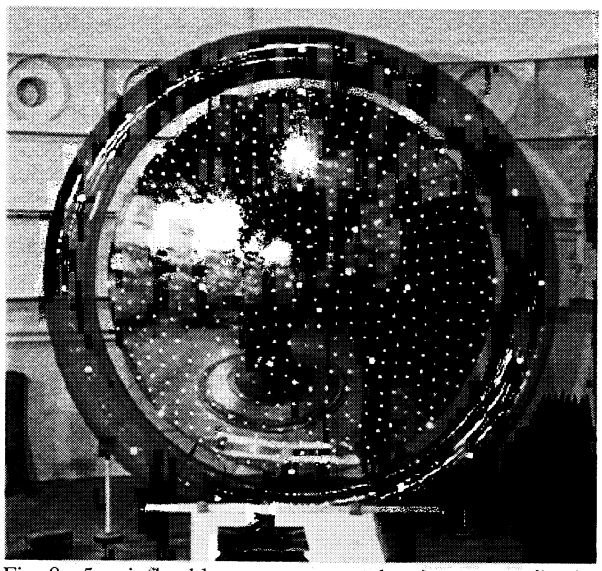

Fig. 8. 5-m inflatable space antenna showing retro-reflective targets illuminated with camera flash.[10]

The antenna's surface profile was measured at more than 500 spatial locations using photogrammetry. Photogrammetry is the process of determining three-dimensional spatial coordinates from two-dimensional photographs.[8] In the last decade, photogrammetry has become a relatively straightforward process with the availability of high-quality digital cameras and a variety of commercial photogrammetry software packages. To determine three-dimensional spatial locations, triangulation calculations are preformed on high-contrast surface features that have been identified and located in two or more of the two-dimensional photographs. For smooth, featureless surfaces, such as the reflective 5-m inflatable space antenna, high contrast features must be artificially added to the structure. Affixed circular retro-reflective targets, visible in Figure 8, provide the necessary contrast when illuminated with a camera flash.

We used a Kodak DCS-760M monochrome 6 mega-pixel camera to acquire the images, and to process them we used *Photomodeler*™ software produced by EOS System Inc. Conveniently, the 5-m antenna has previously been studied photogrammetrically and evaluated as a reflector and solar concentrator.[9,10] This simplified our measurements because the structure already had the necessary retro-reflective targets affixed to its surface. The previous results also provided a useful comparison.

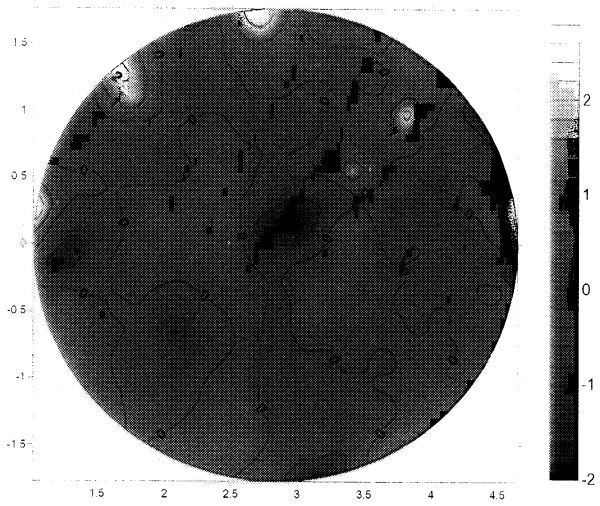

Fig. 9. Contour map of residuals between 6th-order model and measured data. The axes are in meters while the contours in millimeters.

Two batches of photographs were taken, each containing one photograph from eight different locations and viewing angles. A normal photogrammetric measurement requires only one batch of images to determine a single set of x, y, z coordinates for the targets. But we processed two independent batches of images to provide two independent measurement of the location of each retro-reflective target on the structure, thereby providing replicated measurements necessary for *RSM* processing. Since photogrammetry determines the spatial locations in three dimensions for each target, each target had two different unique x and y (and z) locations, one for each batch of images. *RSM* analysis requires replicated points having exactly the same x and y coordinates. To accommodate this need, x and y values were averaged to provide a single x-y location for each target having 2 repeated z measurements.

B. Surface fitting and discussion

We fit several different analytic models to the structure. Three of these models, a "perfect" parabola having the same curvature in x and y, a 6th-order polynomial model, and a piecewise 6th-order polynomial model, are presented and discussed here. In all cases, models were developed for the inner 80% of the parabolic surface only, because the outer limits are not expected to conform to the parabola due to physical manufacturing limitations including tension provided by the cables that attach the antenna to the toroidal outer structure. Additionally, for each fit, five sets of replicated points (10 points in total) were withheld from

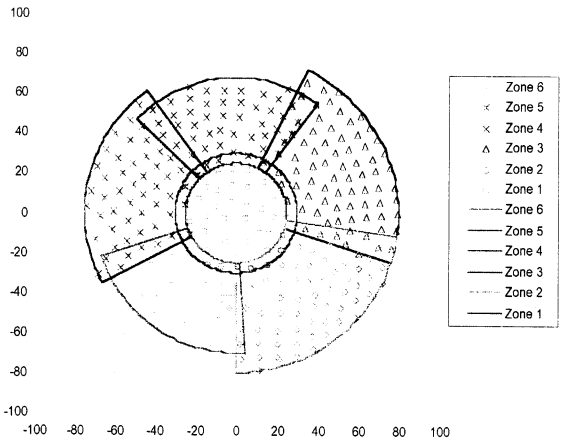

Fig. 10. Segment outlines and measurement points. Axes are in percent of maximum radius.

the fitting process and later used to test the validity of the model.

Initially, an attempt was made to fit the full surface with a sixth-order polynomial using Design Expert by Stat-Ease Inc. Although the surface visually appeared to be a good representation of the structure, a *LOF F-statistic* of 41 indicated that the model was a poor fit. A good fit would be indicated by a *LOF F-statistic* close to unity. This would indicate that the standard deviation of the residuals between the model and the measured data would be similar to the standard deviation of the random measurement scatter at a single point. This is, however, difficult to achieve when the measurement scatter is very low, as is the case with the high precision obtainable from photogrammetric measurements. In this experiment, the standard deviation of difference between the two repeated measurements of z, averaged over all the targets, is 0.26 mm (compared to an overall measured depth of 518 mm over 100% of the antenna's area). A portion of this scatter is due to a 0.09 mm offset in the z direction between the two data sets, but the bulk of the scatter is due to random chance variation.

A plot of the residuals for this model, Figure 9, shows a great deal of structure indicating that the true surface consists of numerous "bumps" with amplitudes approaching 2 mm. Such deviations from the parabolic shape could not be modeled by a single 6th-order polynomial. These result suggested that the surface should either be fit with a higher-order model or that the surface should be partition into smaller regions that might be easier to fit with a 6th-order (or lower) model. The Table-Curve 3D™ program allows higher order response surface models to be fit to the data. However, attempts to fit the full surface with higher order models were unsuccessful: in each case, the *LOF F-statistic* was far from unity and graphs of the residuals exhibited systematic trends similar to those shown in Figure 9.

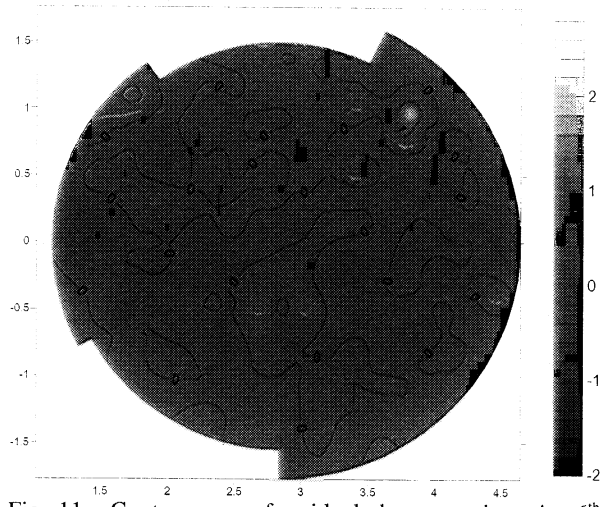

Fig. 11. Contour map of residuals between piecewise 6th-order polynomial model and measured data. The axes are in meters and the contours are in millimeters.

Attempting to fit a smaller area of the surface provided better results. By fitting only the center of the structure, within 30% of the maximum radius, we obtained a *LOF-F statistic* of 1.67 and a *P-statistic* of 0.09, indicating a reasonable fit. The outer area of the structure was then fitted with five additional segments, detailed in Figure 10 and table 1. Each segment overlaps its neighboring segments to reduce the chance of discontinuities forming at the boundaries of the segments. Furthermore, when evaluating the 6-segment model for the surface, a linear blend between the individual fits is performed in the overlap region. This blending function linearly increases the weighing of one segment while decreasing the weighting of the other segment, while maintaining a total weighting of unity. Each segment contained 50 or more targets, with 2 measurements per target location. Note that a full 6th-order model requires a minimum of 28 points. With exactly 28 points, the model would be a perfect fit through every point. However, several of the model terms are insignificant (p is typically 20 for these surface fits) so having 50 spatial points, with 2 samples at each point, allows enough degrees of freedom for the model to generate a realistic fit through the data space. It is possible that if too few measurement locations were used that a wildly oscillating polynomial could be fit to the data. Even if this model producing reasonable statistical parameters such as the *Model F-statistic* and the *LOF F-statistic*, such a model would certainly not predict the test points withheld from the model fit, and consequently would be rejected.

Towards the edge in zones 4, 5 and 6, the antenna deviates significantly from its parabolic shape and the 6th-order polynomial surface fit was unable to follow the curvature even in these segmented regions. Consequently, to obtain a suitable fit, the area had to be further reduced by reducing the radius covered. By creating these slightly smaller regions, we found that it was possible to fit a 6th-order model to each region. As Table 1 indicates, the models pass all the required statistical tests, including successfully predicting data taken at the same time as the rest of the data, but withheld from the model fit. The model predicts the mean shape of the antenna with an average uncertainty of 0.39 mm, with 95% confidence. This precision could be reduced to any desired level of uncertainty by obtaining and processing more batches of images, thereby increasing n_{total} in Eq. (1).

This piecewise approached allowed the development of a model that describes the shape of the surface with a known accuracy. The model can be used to confidently interpolate between measurement points. It could also be used for ray-tracing to find focus spot size and shape. There is an upper frequency limit, however, for which the model is valid for ray tracing. This limit can be found using the Rayleigh smooth-surface criteria. This criteria states that a surface can be consider smooth if

$$\sigma_h < \frac{\lambda}{8\cos\theta} \tag{4}$$

where σ_h is the standard deviation of surface height, λ is wavelength, and θ is the incident angle. Considering the model as a representation of the antenna's average smooth surface, and assuming the residuals (difference between the model and measured points) adequately quantify the surface roughness, the standard deviation of the true surface height σ_h can be calculated from the residuals, resulting in σ_h = 0.39 mm. The assumption that the residuals describe

Segment	Radius	Angle	LOF F Stat	LOF P Stat	Test points passed	*Model Parameters*	$CIHW_{95\%}$ (mm)
Zone 1	< 30%	0 to 360°	1.67	0.09	10/10	23	0.47
Zone 2	25 to 80%	0 to 82°	1.08	0.39	10/10	19	0.22
Zone 3	25 to 80%	72 to 154°	1.20	0.23	10/10	23	0.49
Zone 4	25 to 68%	144 to 226°	1.26	0.25	7/10	25	0.33
Zone 5	25 to 75%	216 to 298°	1.40	0.14	8/10	16	0.43
Zone 6	25 to 70%	288 to 5°	1.39	0.19	8/10	23	0.25

Table 1 – Segment statistics for analysis of parabolic antenna.

the surface roughness is supported by the random nature of and little discernable structure in the residual plot shown in Figure 11. For the worst-case incident angle of θ = 0°, the upper frequencies limit ($f=c/\lambda$, where c is the speed of light) is 95.8 GHz. Although this frequency is well below optical frequencies, the model generated is useful for radio frequencies and could help predict the performance of the antenna for communications and radio-astronomy purposes.

The major advantage of this technique, as opposed to other surface fitting techniques such as spline fits, is the ability to determine the confidence to which the surface is modeled. For example, the *LOF F-statistic* could not be computed using a typical spline fit with only one single set of x, y, z data. The *LOF F-statistic* is a key parameter in determining whether the model accurately represents data, and this is why the approach outlined in this paper includes replicated measurements that allow the computation of this critical statistic for discriminating against poor models. Furthermore, if a given confidence level is required, Eq. (1) can be used to indicate how many measurements are required to obtain that level of confidence. Undoubtedly, spline and related functions are easier to fit to data. In application where visualization of the data is desired, spline fits would be more appropriate than the method outlined in this work. However, if *quantitative* surface shapes are needed than the *RSM* approach can be very useful.

We have successfully produced a model for the majority of the antenna surface. However, the process of developing this piecewise model was lengthy and involved a great deal of trial and error. Several different configurations involving segments of different shapes and sizes were tested before a suitable arrangement of segments, which produced good statistical parameters indicating good fits to the data, was found. It is apparent from our experiences that this process would benefit greatly from the development of an automated system of data analysis. A suggested automated algorithm is:

Fit the full surface with one (6^{th}-order) polynomial. If this polynomial is a bad fit, iteratively reduce the area being fit until a good fit is found. If possible, maintain an area having similar extent in x and y, avoiding long or narrow regions. Repeat as necessary to cover the surface.

It is anticipated that the excluded regions of the antenna, including the outer 20% of the radius, could be accurately fit by further subdividing them into another set of zones. However, without an automated process, this would be very time consuming.

It should be noted that there is a lower limit to the number of measurement points required for a fit, and consequently a lower limit to the size of any one region. If the measurement locations are fixed (as was the case for the 5m antenna), there may be regions of large or complicated curvature where it is not be possible to obtain a model that is a good fit to the data. In this case, not enough sample points have been obtained to clearly identify the shape of the surface. A model can then only be developed if more spatial measurements are obtained, allowing a smaller area to be fitted.

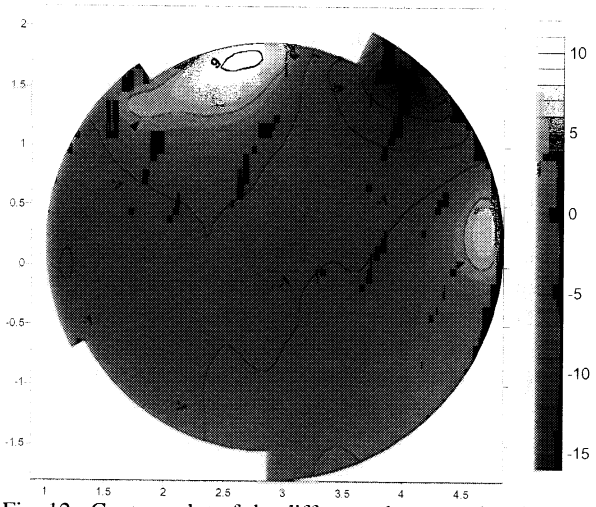

Fig. 12. Contour plot of the difference between the piecewise model and perfect parabola fit. Axes in meters, contours in millimeters.

In addition to the piecewise model, a perfect parabola was fit to the same experimental data to serve as a comparison for the segmented 6^{th}-order polynomial model. Figure 12 shows the deviations of the piecewise 6^{th}-order from a perfect parabola fit to the same data. Note the significant difference at the top and the sides suggesting that the antenna is being vertically compressed, perhaps under the influence of gravity.

From the perfect parabola fit we can also evaluate the focal length of the antenna. We computed a focal length of 118.6 inches, which compares well with previous measurements by Griesch of 120.3 inches on average.[9] This is reasonable agreement considering that the antenna has been deflated and re-inflated between measurements, and that it is currently not in its optimal inflation configuration. Furthermore, Griesch's measurements were made while the antenna was placed in an evacuated chamber.

This study was conducted to see how *RSM* could be applied to improve photogrammetric shape measurements of an inflatable space structure. In the usual *MDOE* approach, measurement requirements are specified at the outset, and an experiment is designed to achieve these goals. One outcome of the present experiment is that this *MDOE* approach can, in fact, be used in a similar manner described in the introduction.

The primary difference between this experiment and the *CARS* experiment is that *CARS* is a single-point measurement technique while photogrammetry

measures hundreds or thousands of points simultaneously. Thus, instead of producing true, independent replicate points, batches of images processed together produce sets of spatial data that are strongly correlated with each other. However, since multiple images from different viewpoints are used in obtaining these spatial points, a degree of randomization has been imposed on the experiment. Furthermore, multiple independently obtained batches of data can be fit to obtain a degree of randomization, while providing replication.

If the photogrammetry data had been less precise, (say 40 times less precise) then the un-segmented 6^{th}-order model would have produced a reasonable fit to the central 80% of the surface, without requiring the partitioning into several spatial regions. This would have allowed much easier implementation of *RSM*, but the additional scatter in the residuals would result in a reduction in the upper frequency limit, determined with the Rayleigh surface roughness criteria, for which the model could be considered valid. Thus, a possible conclusion that can be drawn from this paper is that the approach outlined here is especially good at determining mean surfaces for data (e.g. *CARS*) that has a relatively large amount of scatter compared to the variations in the mean surface. However, if an automated system was developed for partitioning into spatial regions, *RSM* could be implemented in a straightforward manner for quantifying even complex surfaces produced from highly precise data.

Conclusion

Response surface modeling has been applied to two very different optical measurement techniques. *RSM* allowed *CARS* measurements in a scramjet combustor to be performed a factor of 5 times more efficiently than conventional methods. Methods of applying *RSM* to determine the temperature in all three spatial dimensions, or to measure temperature in a plane while varying another parameter were investigated. The minimum number of linear spatial measurement scans that could be used to map a temperature was determined. Finally, the amount of data required to achieve specified levels of uncertainty in derived quantities such as *RMS* temperature fluctuations was computed. In all these cases *RSM* had a pivotal role in the computations.

RSM was also applied to develop a model describing the shape of a parabolic ultra-light weight space antenna using photogrammetry. The shape of the surface over the majority of the antenna was determined with a 95% confidence level uncertainty of 0.4 mm. This uncertainty could be further reduced by acquiring more data. The fitted surface is valid below 95 GHz and could be useful for predicting the performance of the antenna using techniques such as ray tracing.

Acknowledgments

The experiments were carried out at NASA Langley Research Center in Hampton VA. We acknowledge the helpful contributions of Mr. Diego Capriotti, NASA Langley Research Center, for assistance in performing the experiments, and Mr. Richard Pappa and Drs. Phil Drummond at NASA Langley Research Center and Sean O'Bryne, National Research Council Fellow, for helpful discussions.

References

1) Danehy, P. M., DeLoach, R., and Cutler, A.D., "Application of Modern Design of Experiments to CARS Thermometry in a Supersonic Combustor," AIAA Paper 2002-2914, June 2002.
2) Fisher, R. A. "The Design of Experiments," 8th ed. Edinburgh: Oliver and Boyd. (1966).
3) Box, G. E. P., and N. R. Draper, "Empirical Model Building and Response Surfaces," New York, John Wiley and Sons, 1987.
4) See for example, DeLoach, R. "Tactical Defenses Against Systematic Variation in Wind Tunnel Testing," AIAA Paper 2002-0885, 2002; DeLoach, R, J. S. Hill, W. G. Tomek, "Practical Applications of Response Surface Methods in the National Transonic Facility," AIAA Paper 2001-0167, 2001; DeLoach, R, "The Modern Design of Experiments: A Technical and Marketing Framework," AIAA Paper 2000-2691, 2000.
5) Cutler, A. D., P. M. Danehy, R. R. Springer, R. DeLoach, and D. P. Capriotti, "CARS thermometry in a supersonic combustor for CFD code validation," AIAA-2002-0743, 2002.
6) Design-Expert®, Version 6.0.1, Stat-Ease, Inc., Minneapolis, MN, http://www.statease.com.
7) TableCurve®3D User's Manual, Version 3.0, © by AISN Software, Inc., 1997.
8) K.B. Atkinson, *Close range photogrammetry and machine vision*, Whittles publishing, Scotland, 1996.
9) L.R. Giersch, *Pathfinder photogrammetry research for ultra-lightweight and inflatable space structures*, NASA/CR-2001-211244, 2001.
10) R.S. Pappa, L.R. Giersch, and J.M. Quagliaroli, *Photogrammetry of a 5m inflatable space antenna with consumer digital cameras*, NASA/TM-2000-210627, 2000.